

# **FINAL REPORT**

## **FOR VIRTUAL SCAVENGER HUNT PROJECT**

### **DEVELOPED FOR AOARD BY NOVAMENTE LLC**

**March 20, 2009**

## **1 Summary**

The Virtual Scavenger Hunt (VSH) project was carried out by Novamente LLC for AOARD during June 2008 – February 2009. It involved the development of an environment within the Multiverse virtual world, oriented toward allowing individuals to acquire and reinforce skills via interacting with each other and with AI-controlled agents in a class of virtual-world based gaming situations called ***scavenger hunt scenarios*** (scenarios in which teams combining humans and AIs compete to find hidden objects in complex environments). The family of scavenger hunt scenarios was chosen because of its value in emulating various relevant real-world situations; and also because of its amenability toward allowing the study of various aspects of human behavior, interaction and communication in the context of game-play.

The project had two phases:

1. The most extensive phase: creation of the Virtual Scavenger Hunt game, along with supporting software for game-playing and instrumentation.
2. An experimentation phase during which humans interacted with the game, with a view toward identifying properties that would be interesting to explore in future crosscultural studies

Goals for phase 1 were:

- **G1:** Creation of the AvatarBrain software for controlling AI agents, building on the foundation of Novamente's PetBrain software
- **G2:** Creation of a scavenger hunt scenario in the Multiverse virtual world, in which humans and AIs can collaboratively play scavenger hunt games

Goals for phase 2 were:

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE <b>11 FEB 2010</b>		2. REPORT TYPE <b>Final</b>		3. DATES COVERED <b>17-06-2008 to 17-09-2009</b>	
4. TITLE AND SUBTITLE <b>Virtual scavenger hunt: an AI-powered virtual environment designed for training individuals in effective teamwork, and analyzing cross-cultural behavior</b>			5a. CONTRACT NUMBER <b>FA48690814105</b>		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) <b>Benjamin Goertzel</b>			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) <b>Novamente ,LLC,1405 Bernerd Place,Rockville,MD,20851</b>			8. PERFORMING ORGANIZATION REPORT NUMBER <b>N/A</b>		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) <b>Asian Office of Aerospace Research &amp; Development, (AOARD), Unit 45002, APO, AP, 96338-5002</b>			10. SPONSOR/MONITOR'S ACRONYM(S) <b>AOARD</b>		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) <b>AOARD-084105</b>		
12. DISTRIBUTION/AVAILABILITY STATEMENT <b>Approved for public release; distribution unlimited</b>					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT <b>The Virtual Scavenger Hunt (VSH) project was carried out by Novamente LLC for AOARD during June 2008 ? February 2009. It involved the development of an environment within the Multiverse virtual world, oriented toward allowing individuals to acquire and reinforce skills via interacting with each other and with AI-controlled agents in a class of virtual-world based gaming situations called scavenger hunt scenarios (scenarios in which teams combining humans and AIs compete to find hidden objects in complex environments). The family of scavenger hunt scenarios was chosen because of its value in emulating various relevant real-world situations; and also because of its amenability toward allowing the study of various aspects of human behavior, interaction and communication in the context of game-play.</b>					
15. SUBJECT TERMS <b>sociocultural modeling, Virtual Reality, human performance optimization</b>					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT <b>Same as Report (SAR)</b>	18. NUMBER OF PAGES <b>11</b>	19a. NAME OF RESPONSIBLE PERSON
a. REPORT <b>unclassified</b>	b. ABSTRACT <b>unclassified</b>	c. THIS PAGE <b>unclassified</b>			

- **G3:** Explore methods of measuring AI, human and collective game playing and learning ability
- **G4:** Explore methods of studying cross-cultural differences in game playing performance, with a focus on understanding the individual and group traits that make teams more or less effective in different cultural contexts

The phase 1 goals were met successfully, although some scalability problems were met with the Multiverse platform, which (unless further software work is done) would prevent large-scale, multiplayer scavenger hunt games from being played. Nevertheless, a great deal can be learned from games with a relatively small number of players; and these scalability problems can be surmounted with a modest amount of further engineering work.

Regarding the phase 2 goals, we focused specifically on gameplay in which AIs search for hidden objects, under the (loose or rigid) control of humans giving them commands. In this context, we concluded that measuring the relative speed of solution of scavenger hunt problems is a simple but effective way of assessing AI capability, so long as one looks at speed of solution in both AI-only and AI/human combinations. We also identified several dimensions that may be interesting to explore in future crosscultural studies:

- Coarse statistics of human/AI interaction: we found preliminary, suggestive evidence that ***individuals from different cultures might differ in the frequency of commands they give to AIs***, as well as to how markedly they change this frequency in reaction to AI success/failure
- Communication about AI actions: in particular, we found preliminary evidence that ***individuals from different cultures might have differential tendency to seek explanations of AI motivations***.
- Implicit communication between human players, in scenarios where multiple human players are on a single team, but textual and verbal communication between human players is not enabled: we found preliminary suggestive evidence that ***individuals from different cultures might have differential capability in implicit communication***

A paper is being prepared for publication regarding the personality modeling aspect of the work done for G1, in which special logical rules were created to govern the behaviors of the AIs taking place in the scavenger hunt game. Preparation of a paper regarding the crosscultural, comparative study aspect will need to wait until detailed studies of human behavior can be done, further exploring the preliminary, suggestive observations made above.

## 2 Detailed Report

Here we detail the progress that was made toward the four goals described above.

### G1: Creation of AvatarBrain for Controlling AI Agents

During 2007-2008, Novamente LLC has developed a software system called the PetBrain, a proprietary AI system based on the Novamente Cognitive Engine architecture. The PetBrain controls smart pets in virtual worlds. In the course of the VSH project, we have modified it to control humanoid agents, creating a variant of the PetBrain called the AvatarBrain.

Virtual humanoids, running in the AvatarBrain, share the same basic cognitive architecture used for pets in the PetBrain. The PetBrain was modified to support different agent types (pets, humanoids, etc). Pragmatically, the user may input the agent's type into the Multiverse console using a `/loadagent` command, e.g.

```
/loadagent Fido dog
```

```
/loadagent Clara humanoid
```

We also improved the vision sensing in the PetBrain, in the course of creating the AvatarBrain, enabling agents to have a realistic sense of what lies in their field of view. The AvatarBrain vision sensing system uses a technique called frustum culling to select all the objects that are inside the agent's Field of view. Besides selecting the objects that are inside agent's FOV, this sensor must remove from the agent's memory the hidden objects (objects positioned behind other objects that are not visible by the agent)

The following commands were implemented to be used by the human player, within the Multiverse console, during a Scavenger Hunt match

- `go find treasures [agent name]:` make the agent look for treasures
- `follow me <agent name>:` make the agent follow the owner indefinitely
- `follow me collect <agent name>:` make the agent follow the owner indefinitely, but if a treasure was seen by the agent, it will collect that treasure. After fetch the treasure to it's base, the agent will go after it's owner to follow him/her again.
- `come here <agent name>:` make the agent come near to it's owner and after a short period it will go back to it's normal activities

- `explore area <area number 0-8> <agent name>`: make the agent go to a specific area to explore. If that area was already explored by the agent, the command will be ignored. If a treasure was discovered during the exploring process, the agent will fetch that treasure to its base.
- `wait <agent name>`: The agent will stay in its current position indefinitely until the owner gives another command.
- `regroup team <team name blue|red>`: Make all the agents of a given team to go back to their base
- `spy <spied agent name> <agent name>`: Make an agent look for another. If the spied agent wasn't seen or was lost by the spy agent, the latter one will turn around itself to find the target. Given that the spied agent was identified the spy will get closer to it.
- `i blue team | i red team`: These are special commands which can be used by human avatars to become member of a specific team (red or blue). After the owner of the humanoids gives the command "lets play scavenger hunt", there is a short period which the agents organize the team members. During that period, the human avatars can use one of these commands to become a team member. After that period these commands have no effect. The finish of the mentioned period is determined by the moment of the agents go to its bases and wait for the game start command "go find treasures".

An example server log resulting from a human interacting with an AI agent using this syntax is given below. The lines containing "instruction" reflect instructions typed by the user to the AI agent. The other lines reflect actions taken by the AI agent based on commands from the AvatarBrain. The "..."'s occur where we have admitted lengthy series of "walk" commands issued by the AvatarBrain.

```
avatar Paula instruction: lets play scavenger hunt
avatar Saulo instruction: i red team
avatar Paula instruction: i blue team
avatar 'Saulo' walk to position(x=-75846.0,y=216910.0,z=0)
...
avatar 'Saulo' walk to position(x=-81045.0,y=208526.0,z=0)
avatar Paula instruction: go find treasures
avatar 'Saulo' walk to position(x=-81362.0,y=209310.0,z=0)
```

```
...
avatar 'Paula' walk to position(x=-101005.0,y=231653.0,z=0)
avatar 'Saulo' grab accessory(759250)
avatar 'Saulo' walk to position(x=-108938.0,y=242448.0,z=0)
...
avatar 'Saulo' walk to position(x=-85903.0,y=219811.0,z=0)
avatar 'Paula' grab accessory(759248)
avatar 'Saulo' walk to position(x=-120925.0,y=244924.0,z=0)
...
avatar 'Saulo' walk to position(x=-81625.0,y=229875.0,z=0)
avatar 'Paula' drop
avatar 'Saulo' drop
avatar 'Saulo' walk to position(x=-80541.0,y=209655.0,z=0)
...
avatar 'Saulo' walk to position(x=-110071.0,y=206941.0,z=0)
avatar 'Saulo' grab accessory(759245)
avatar 'Paula' walk to position(x=-94671.0,y=216256.0,z=0)
...
avatar 'Saulo' walk to position(x=-92582.0,y=209211.0,z=0)
avatar 'Paula' grab accessory(759243)
avatar 'Saulo' walk to position(x=-92157.0,y=209174.0,z=0)
...
avatar 'Paula' walk to position(x=-94640.0,y=220269.0,z=0)
avatar 'Saulo' drop
avatar 'Paula' walk to position(x=-94208.0,y=220677.0,z=0)
...
avatar 'Saulo' walk to position(x=-115204.0,y=213349.0,z=0)
avatar 'Paula' drop
avatar 'Saulo' walk to position(x=-115191.0,y=214753.0,z=0)
...
avatar 'Saulo' walk to position(x=-110375.0,y=225375.0,z=0)
avatar Paula instruction: save_map
avatar Paula instruction: save_vismap
avatar 'Saulo' walk to position(x=-97570.0,y=227334.0,z=0)
```

## **G2: Adaptation of Virtual World Software for Scavenger Hunt Games**

Using the Multiverse platform, we created a “serious game” scenario analogous to a scavenger hunt. Multiverse’s technology allows one to create multiplayer games in virtual worlds with a relatively small time and cost investment. Novamente’s virtual pets were previously integrated with an existing simple virtual world within Multiverse.

The high-level goals for the scenario were:

1. Allow exploration of both cooperative and competitive interaction between human-controlled avatars
2. Allow human-controlled and AI-controlled avatars to interact as opponents and as teammates
3. Provide an opportunity to train AI agents to imitate human players and study their imitative accuracy
4. Provide an avenue for studying and comparing the social, strategic and tactical behavior of individuals from different cultures (and the training of AI’s, via imitation, to enact roles in the game in a manner similar to individuals from various specific cultures)

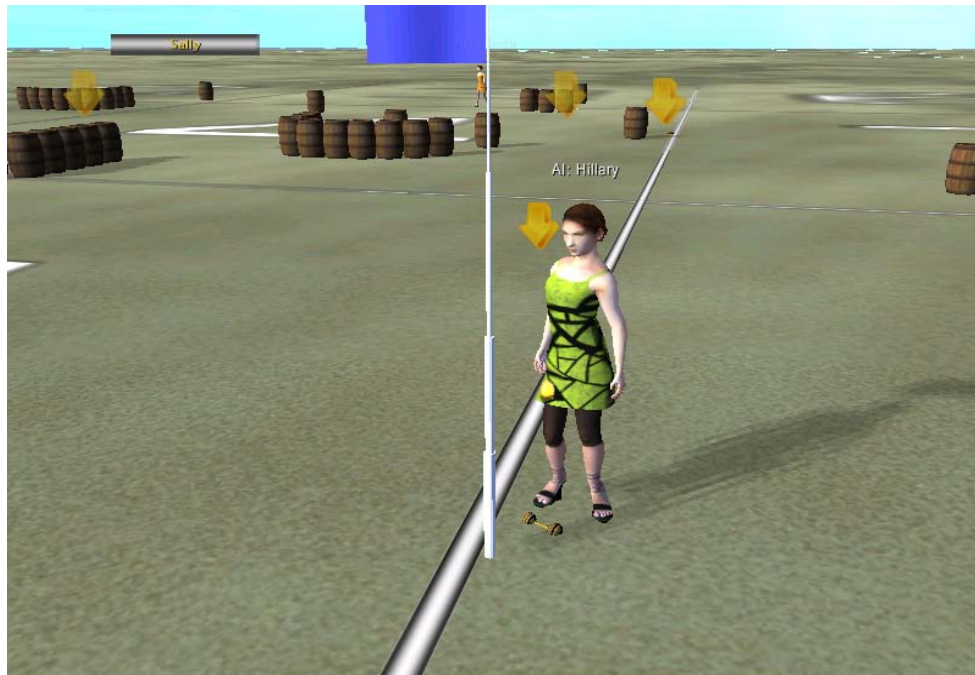
We succeeded at achieving goals 1, 2 and 4 in the course of the work, but did not end up having time to explore 3 except in very simplistic ways. However, the framework we created is ideal for the exploration of 3 so we hope to do so in future work.

In the scavenger hunt game as implemented, a number of teams compete to achieve the objective of finding hidden objects in the virtual world. Teams may consist of human-controlled avatars, AI avatars, or a combination of the two. In general, the game may be played either in a specific environment constructed just for the game, or more generically across large regions of the virtual world. However, all our experiments took place in a specific scavenger hunt environment. And we focused specifically on gameplay in which AIs hunt for hidden objects, but under the supervision of humans who give them commands.

Operationally, what was created was an online multiplayer scavenger hunt game, including:

- A downloadable client to be used by humans;
- A proxy that allows integration with the AvatarBrain, thus enabling AI players;
- A virtual world with an environment embodying a scavenger hunt scenario, and support for persistent avatars.

Following are several screenshots of our initial test scenario and some AI and human agents playing scavenger hunt in it. The human controlled agents are the ones wearing halter tops.







In the specific scavenger hunt game we implemented, a human avatar hides an object and one NPC tries to find the treasure, based on the human's (loose or rigid) instructions. Future implementations will incorporate the possibility of creating teams of NPCs that will cooperate among themselves to find one or more hidden treasures.

To start a match of Scavenger Hunt, the human player must load some humanoids into the world using commands such as (typed into the Multiverse client console):

```
/loadagent humanoid Maria
```

Then, the player must to give the command: "lets play scavenger hunt" to start playing. During a short period after that command, the agents will go to their team bases and wait for the game start command. After that, the human player must to give the command: "go find treasures" to start the game.

### **G3: Measure the AI's Game Playing and Learning Ability**

One important data analysis problem raised by the VSH project is: What makes a player, or a team, effective? We did not carry out a sufficient number of experiments to thoroughly explore this sort of data analysis. However, we did make some simple observations, using our default scavenger hunt scenario in which there are 8 hidden objects in an area of approximately 1400 sq ft, and using a gameplay mode in which AIs find hidden objects under human guidance.

In this context,

- Unaided agents take ~10min on average to complete the scenario.
- Aided agents take ~7min
- Humans take around 4min.

So, our current AIs are not as smart as humans at solving the scavenger hunt – but getting human instructions helps them. One meaningful way to compare humans, AIs and human/AI combinations is via how close AI performance comes to human performance (and using appropriate AI algorithms it might eventually supersede human performance, of course).

### **G4: Study Cross-Cultural Differences in Game Playing Performance**

We believe that the scavenger hunt scenario will provide an excellent context in which to compare human individual and team behavior, as it depends on various aspects including the cultural and educational backgrounds of the

players, the size of the teams, the specific nature of the game scenario and so forth. Hypothetically, one could set up a machine learning problem whose specific goal was to learn patterns distinguishing the gameplay of individuals from different cultures; and human inspection of these learned patterns may be extremely informative. Furthermore, if one trained AI avatars to imitate the behavior of players from different cultures, one could then study the aspects of players from different cultures that are hardest for AI's to emulate – which is an interesting way of honing in on the subtlest cultural differences.

In this small project we could not explore these issues too thoroughly, but we made some simple, heuristic observations regarding crosscultural differences, due to the fact that our development team was international, spanning the US, Brazil, and China.

For example, we found preliminary, suggestive evidence that ***individuals from different cultures might differ in the frequency of commands they give to AIs***, as well as possibly differing in how markedly they change this frequency in reaction to AI success/failure. Some individuals command the AIs frequently, so much so that the AIs don't get the chance to complete one command before the next one is issued. Others tend to allow the AIs more room to explore on their own. We noted that a number of our American staff members displayed a strong tendency to allow AIs freedom to explore, as compared to individuals of other nationalities.

On the other hand, what happens when an AI that was given freedom fails to find an object, even though its location was obvious? Does the human player react by micromanaging the AI the next time, or does he continue to give the AI freedom? Our preliminary observations on this were less clear and we would like to do more in-depth studies.

Next, we found hints of significant cultural differences regarding communication about AI actions: in particular, we found preliminary evidence that ***individuals from different cultures might have differential tendency to seek explanations of AI motivations***. Some individuals described what AIs are doing using languages like “he wanted ...”, “he thought...” and others did less so. We found some suggestive evidence that our staff from non-Western cultures might have more propensity to anthropomorphize the AI agents; but to validate this sort of observation would require more careful treatment of crosscultural linguistic issues, because all of our staff were communicating in the context of the game in English, regardless of their first language.

Finally, while observing our staff play the game in situations with multiple human players and a single AI player on a team, it occurred to us that it might be interesting to study *implicit communication* between human players. Toward this end we did a few experiments with gameplay in which multiple human players are on a single team, but textual and verbal communication between human players is not enabled. Through this experimentation, we found preliminary suggestive evidence that ***individuals from different cultures might have***

***differential capability in implicit communication.*** Unsurprisingly, we found that teams consisting of individuals with a history of working together in other contexts, appeared better at implicit communication; but we also found intriguing hints that our Chinese staff displayed a greater knack for implicit communication than our other staff. However, we stress that these are extremely preliminary, rough observations that would require much more extensive experimentation and careful validation before they could be considered conclusive in any way.